Massive Scientific Visualization

Philip D. Heermann
Sandia National Laboratories
Albuquerque, New Mexico
pdheerm@sandia.gov

Sandia National Laboratories

Virtual Experiments

Simulation of Richtmyer-Meshkov Instability.

Two gases, which are initially separated by membrane pushed against wire mesh, are subjected to Mach 1.5 shock.

ASCI SST Machine
960 nodes.

Visualization by M. Duchaineau, J. Hobson, D. Schikore, LLNL

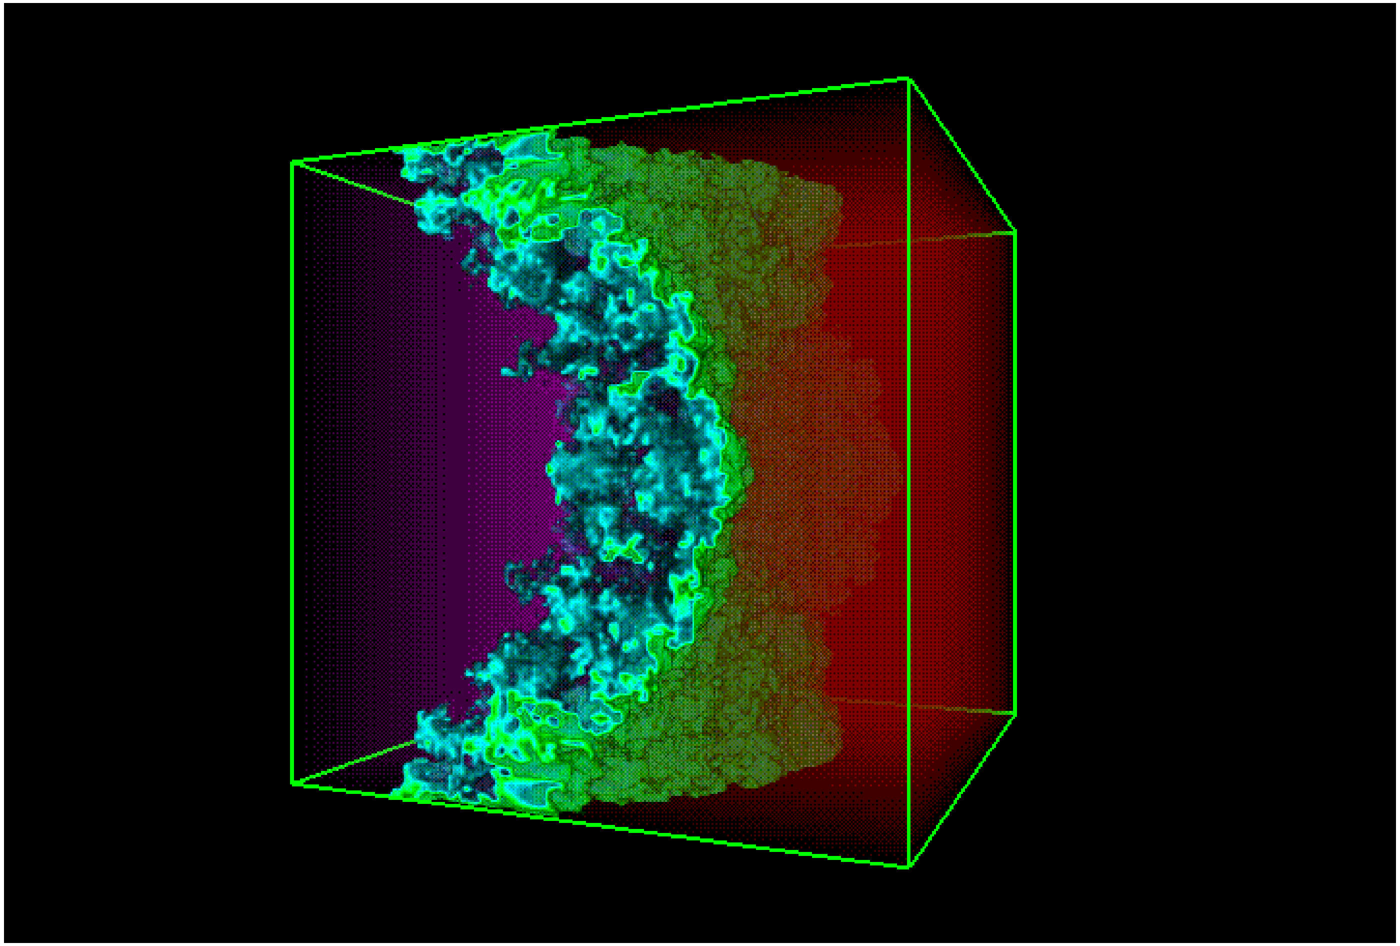

Sandia National Laboratories

Virtual Experiments

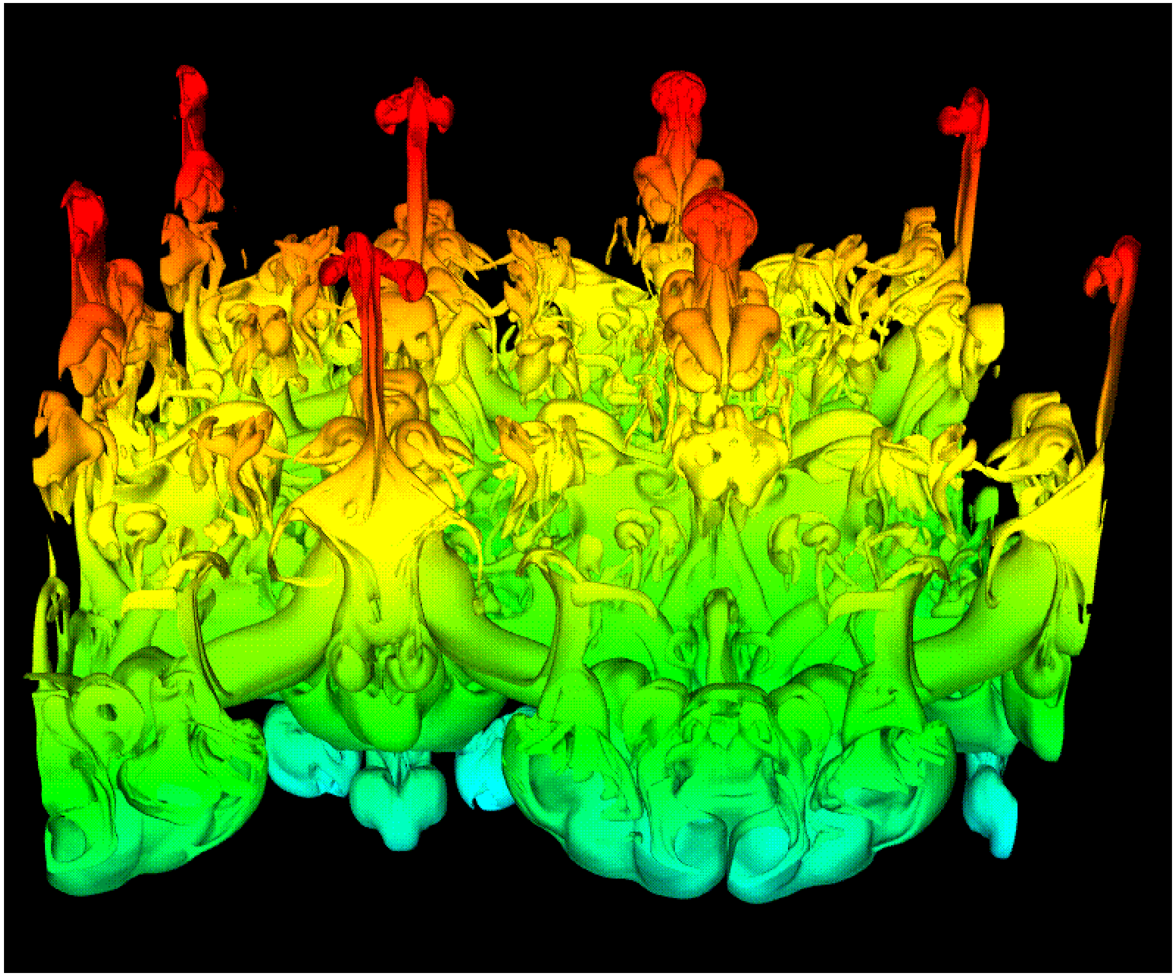

Very late time non-linear development of an unstable shocked interface between two fluids. The density isosurface shown was extracted and rendered from a 73.5 million cell AMR mesh.

Image: Bob Kares/Jamie Painter, LANL

Virtual Experiments

Shock Physics Calculation:

Re-entry Body Impact

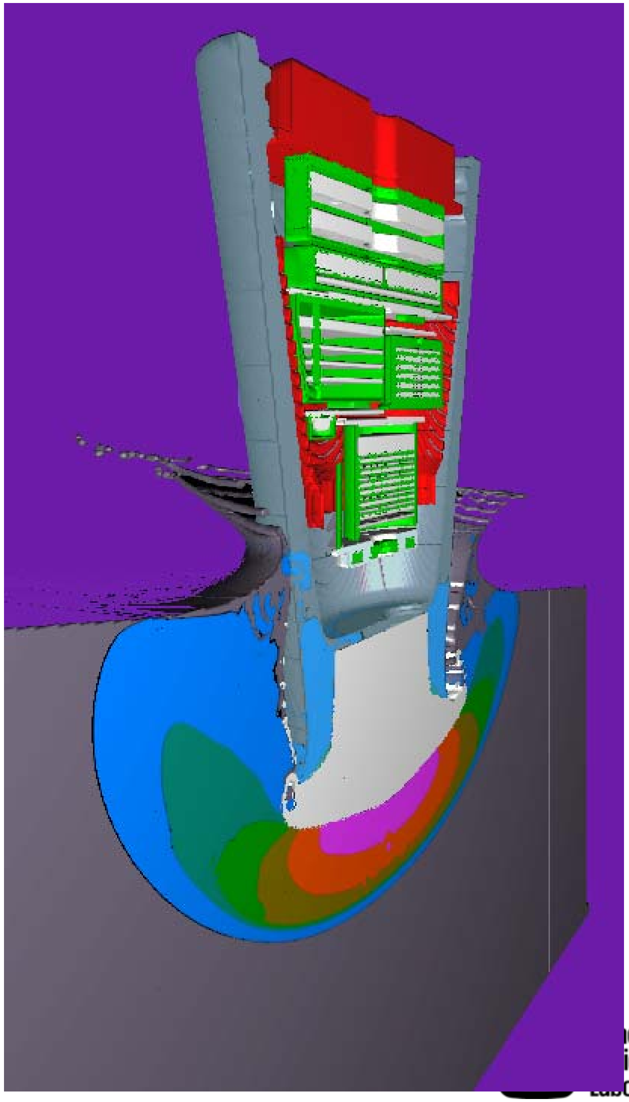

Some Data Facts

- Recent ASCI computation (LANL on LLNL White machine)
  - 468 million cells (max), 14-25 variables, 364 timedumps, 21 TB of graphics files (652 TB total archived data, including restart files)
  - Ratio of extracted geometry to total mesh geometry nominally around 4-5% (per empirical observation)
  - One example: 355 M cells (18.86 GB), extracted geometry 31.6 M triangles (773 MB)
- PPM SC99 Gordon Bell Data Set
  - 8 billion cells, 2 variables, 273 timedumps, ~4.3 TB, ~300 GB compressed
  - One variable: ~ 8 GB, ~550 MB compressed
  - One particular extracted isosurface: ~470 M polygons, ~5 GB compressed

•ASCI is dealing with very large data requirements

•Extracting geometry is one effective way to reduce data before moving

•Even when data is big, it is often possible to comfortably fit one or more full-size data objects on a desktop/laptop local disk

•Albeit somewhat pathological, the size of extracted geometry can substantially exceed that of its corresponding raw data

Sandia National Laboratories

VIEWS Functional Architecture

Data Sources: Simulations, Archives, Experiments, Sensors

Data Services:
- Permutation $M \Rightarrow N$
- Filtering
- 1D/2D Subsetting
- Data Algebra $x,y,z \Rightarrow mag/\Phi$
- Format/Representation Conversion
- Data Reduction
- Data Serving

Information Services:
- Feature Detection and Extraction
- Data Fusion & Comparison
- Visual Representations Generation (eg. isosurfaces)
- Volume Visualization Preparation (eg. opacity assignment, resampling ...)

Visualization Services:
- Surface rendering
- Volume rendering
- Runtime services
- Multi-Visualization Technique Combine
- Time Sequence Generation

Display Modalities:
- Desktop Display
- Theater Display
- Powerwalls
- Immersive Stereoscopic

Users Services:
- Navigation
- Rendering Control
- Advanced User Interface
- Collaborative Control
- Display Control

The Big Picture

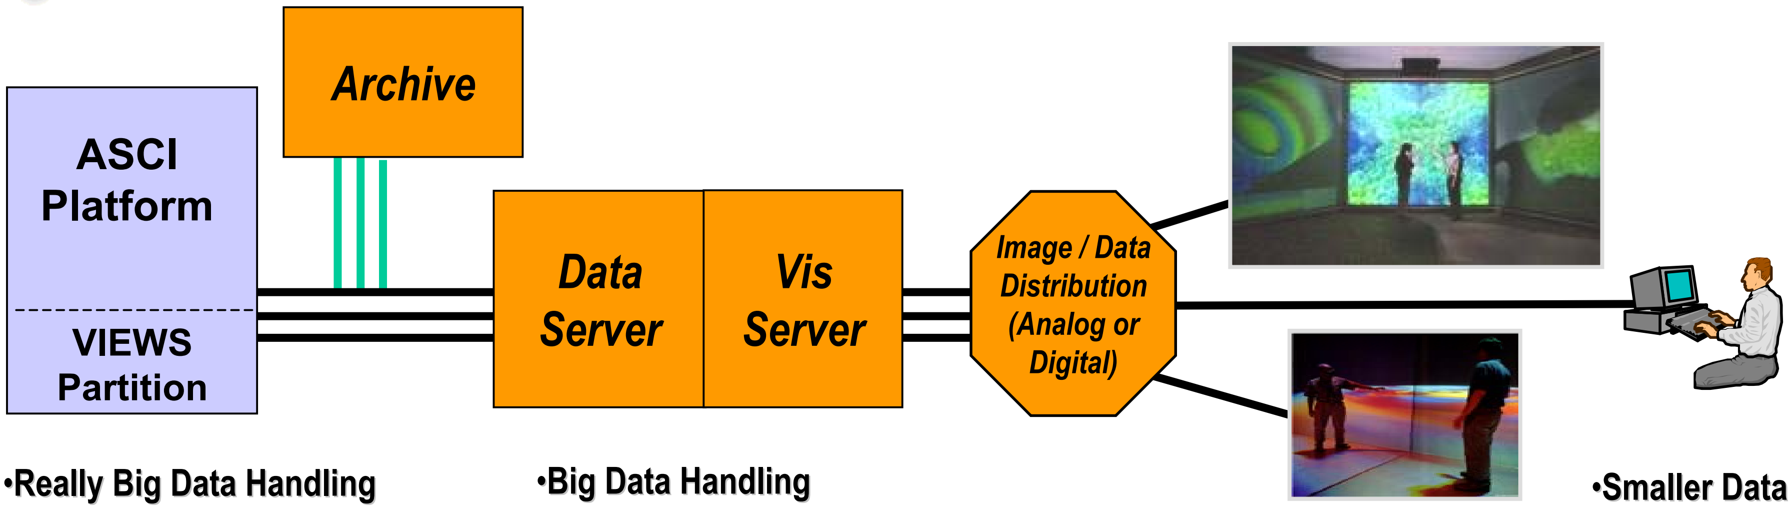

- Really Big Data Handling
- Data Manipulation
  - Data Subsetting
  - Geometry Extraction
  - Feature Detection / Extraction (Data Mining)
  - Data Preparation for
    - Hierarchical / Multi-res
    - Out-of-core
- Some Full-featured Vis and Rendering (maybe SW only)

- *EnSight Server of Servers*
- *VisIt / Parallel VTK*

- Big Data Handling
- Data Manipulation
- Full-featured Visualization
  - Raw data
  - Hierarchical / Multi-res
  - Out-of-core
- Parallel HW Rendering

- *EnSight Server of Servers*
- *EnSight Client*
- *VisIt / Parallel VTK*
- *TeraScale Browser*

- Smaller Data
- Remote Image Display
- Visualization
  - Hierarchical
  - Out-of-core
  - Anything with smaller data
- HW Rendering

- *EnSight Client*
- *EnSight*
- *VTK*
- *TeraScale Browser*

Sandia National Laboratories

We are developing approaches that support rendering to high pixel-count and/or tiled displays

Sandia/California's Power Wall
4x3 tile-display
Approximately 16 Mpixels

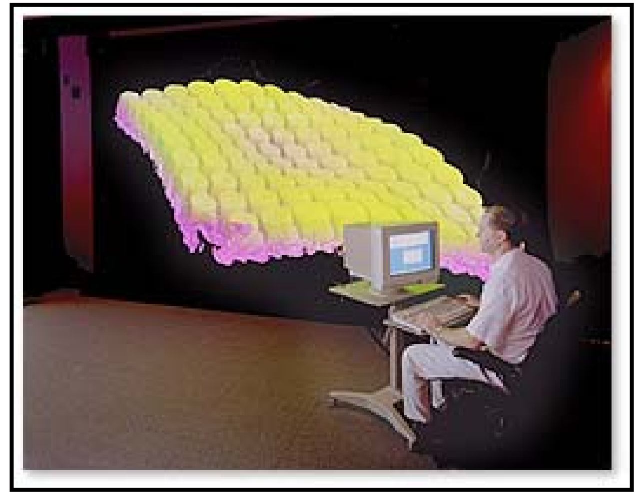

Sandia/New Mexico Facility
Deployed for 48-tile, 3-screens (60 Mpixels)

Rendering the 470 million-triangle data to 20 Mpixels: ~80 million triangles/sec, ~6 sec per frame

IBM's
21-inch
9 Mpixel
desktop
display

62 Million Pixels driven by commodity PC cluster

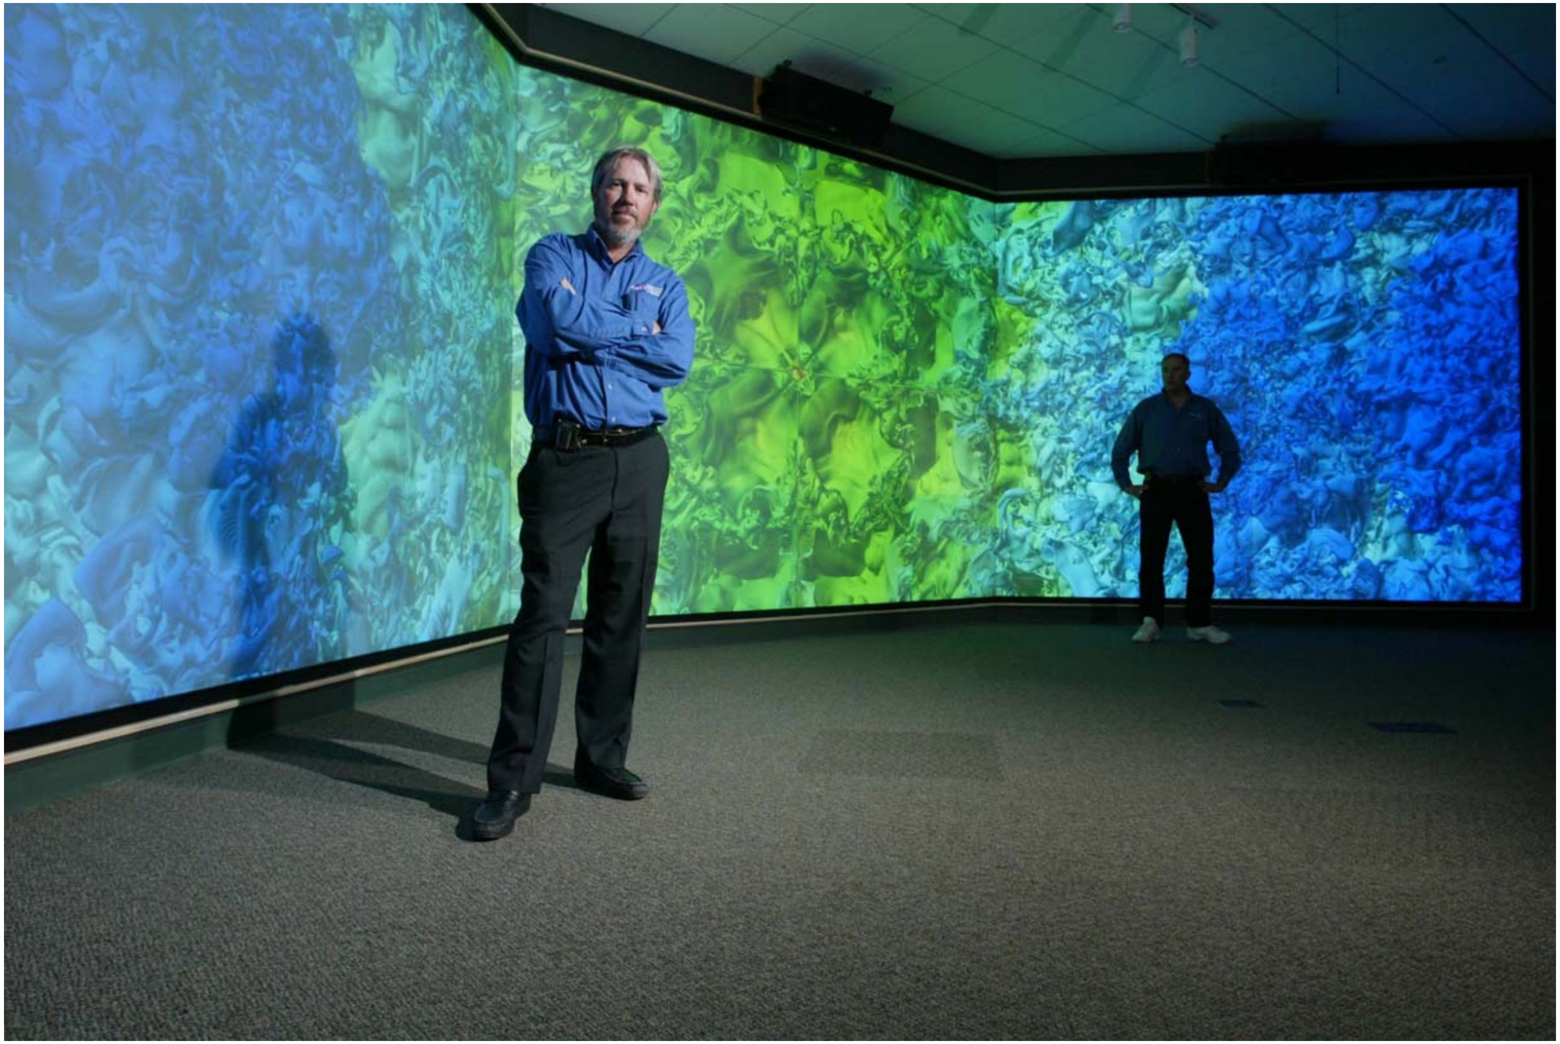

At 15 fps, the system can display 2.7 GB/sec

ASCI/VIEWS is working with partners to develop scalable end-to-end solutions for large data visualization

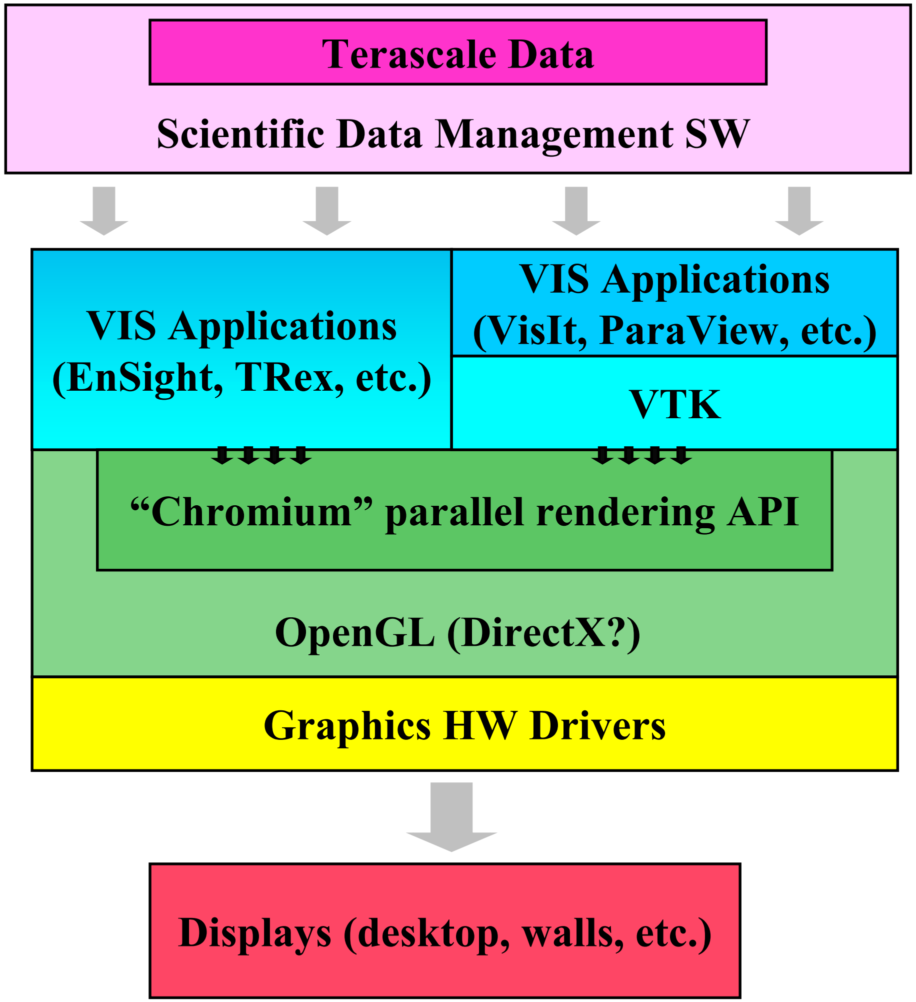

ASCI Tri-labs:

- Sandia National Laboratories
- Lawrence Livermore National Laboratory
- Los Alamos National Laboratory

Partners (past and present):

- CEI (Parallel EnSight)
- Kitware (Parallel VTK)
- Stanford University (see "The Chromium Project" at sourceforge.net)
- Princeton University (Scalable Displays)
- NVIDIA (Linux graphics drivers)
- RedHat (Chromium and Distributed X Server)
- IBM (Bertha displays)

Sandia National Laboratories

Price-performance of commodity-based graphics clusters is very attractive …e.g. Sandia's Europa Cluster

- 128 nodes (Dell Precision 530 workstations)
- Each with dual 2.0 GHz Pentium-4 Xeon CPUs, 1 GB RDRAM, GeForce3 (NVIDIA) graphics cards
- Myrinet 2000 interconnect
- Linux and Windows operating systems

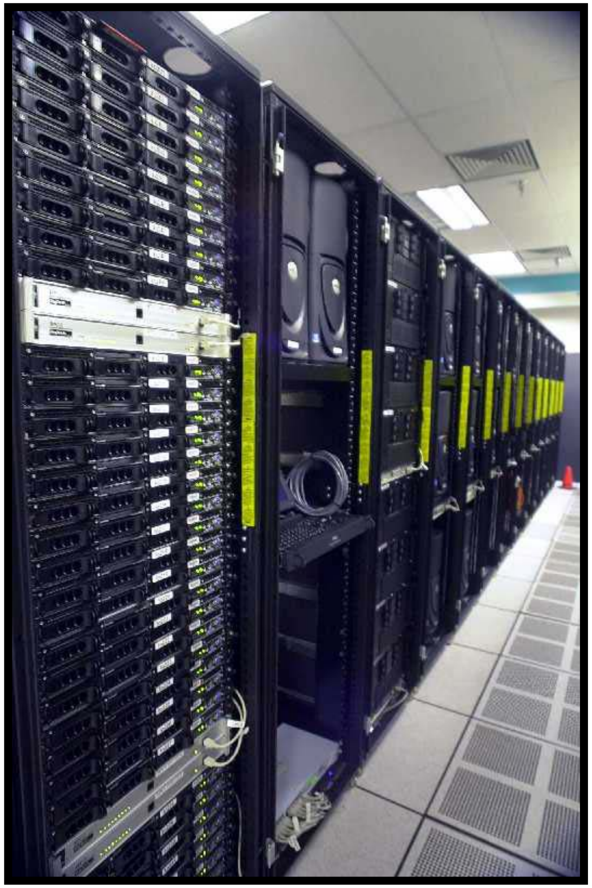

- Aggregate CPU power: 0.5 teraflops (peak), ranked 103 on November, 2001 Top500 list
- Aggregate GPU power: 100 teraops (peak), ~3 Billion polygons per second (peak) **Demonstrated rendering at 1 Billion polygons per second**

- **Total cost: ~$900K**

Sandia National Laboratories

We are benefiting from the Game market's impact on PC-graphics … could games boxes be in our future?

Sony PlayStation:

- Floating Point Performance 6.2 GFLOPS
- DRAM Bus bandwidth 48GB per Second
  - (Bus width 2560 bits)
- Maximum Polygon Rate
  - 66 Million Polygons per Second

And the X-Box:

- 800 GOP at ~$200
- Maximum Polygon Rate
  - 150 Million Polygons per Second

- **Complex Software systems may not fit!**

Miscellaneous Observations

- Problem setup is a huge bottleneck
  - Can take 6-12 months to build a big/complex mesh
- Big simulations run for a week or more
  - Usually means post processing visualization
- More polygons than pixels becoming common.
  - Transition in Large Data Visualization ?
- Rendering diminishing as major bottleneck
- Data handling consumes more wall clock time than visualization -> data services
- More selective archival needed
  - tapes are too small
- No steering of big computations; yes tracking
  - Reproducibility of results important
- Scalable rendering
  - hybrid system is likely the best (combine sort first & last)

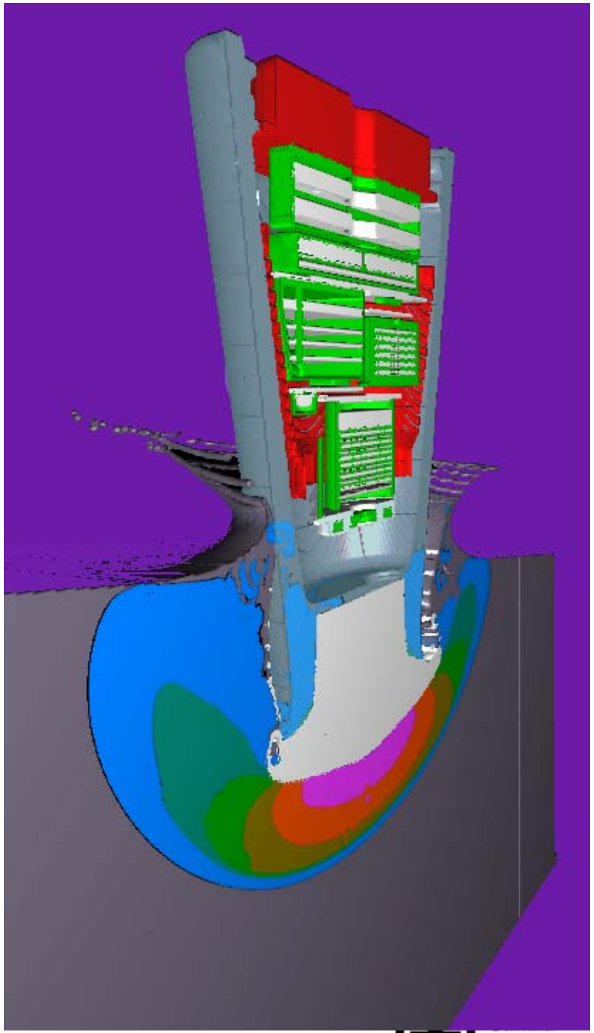

National Laboratories